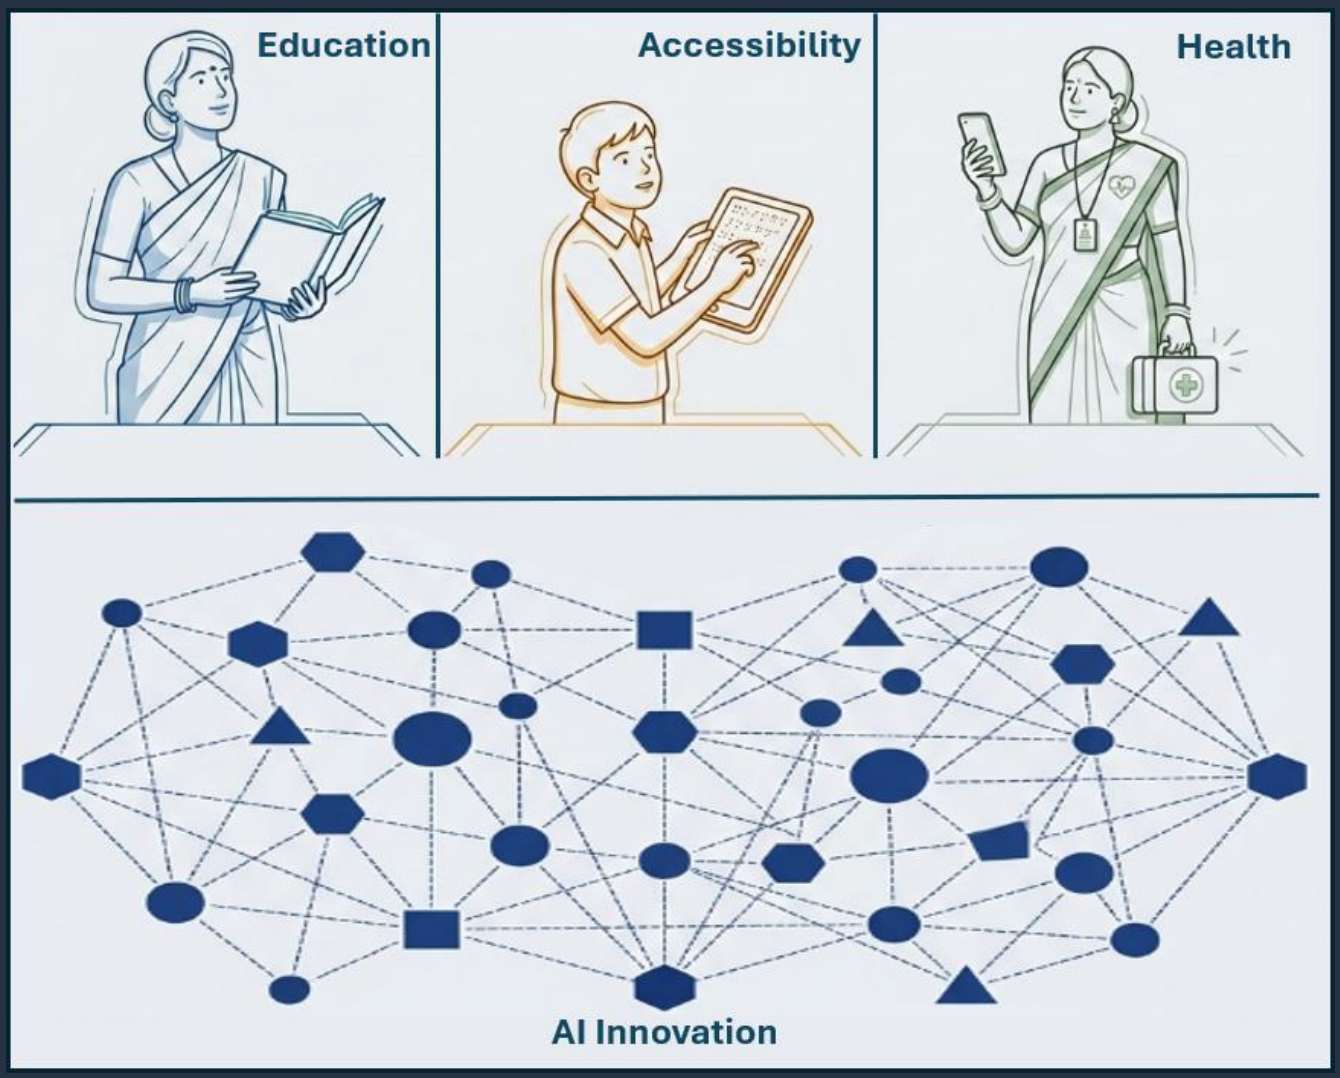
AI INNOVATION & INCLUSION INITIATIVE

A Microsoft India and IIIT-Bangalore Initiative

A4I Annual Impact Report 2025- 2026



Draft for review

Contents

About A4I and its first year journey 3

✦ From Vision to Reality 5

↗ From Pilot to Scale 5

❖ The A4I Ecosystem 6

▦ Year One at a Glance 8

✎ Sikshana Foundation.....11

◉ Vision Empower.....14

✚ Khushi Baby 17

⚙ A4I Engineering.....21

◆ A4I Research 23

★ What We Learned 26

➜ The Path Ahead 27

Closing Note 27

About A4I and its first year journey

AI innovation, translated for public value

A4I was created to solve a stubborn problem in social innovation: promising AI solutions demonstrate potential in smaller pilots, but struggle to become dependable tools at a large scale in the hands of teachers, frontline workers, and inclusion-focused educators.

With the support and collaboration of Microsoft India, A4I was launched in March 2025 with the belief that advanced AI should not remain limited to well-resourced settings. Its value must be tested in the environments where the need is greatest and where public systems, frontline workers, teachers, and communities operate under real constraints.

In Year 1, A4I brought together IIIT-Bangalore, Microsoft India, and three domain partners across education, accessibility, and healthcare. The initiative is built on a simple belief: population-scale impact is not an "AI problem" alone, it is an ecosystem challenge. Industry brings frontier capabilities, academia brings research rigor and scale engineering, civil society anchors relevance and adoption, and government provides the pathway to legitimacy, delivery infrastructure, and statewide reach. A4I serves as the translation layer that turns promising AI assets into usable, safe, and continuously improving public-good solutions.



The initiative supported field deployments, strengthened partner capability, and converted early evidence into reusable engineering and research practices. Year 1 established A4I as a live translation hub, working through the difficult middle layer between promising pilots and population-scale public value. This layer includes engineering reliability, contextual design, partner enablement, field research, and responsible governance.

This report presents Year 1 as a foundation year. It not only records reach and early adoption, but also documents how the last-mile impact was achieved. It also recognizes the difficult lessons that emerged from the field: the importance of workflow fit, the need for human oversight, the role of trusted intermediaries, and the responsibility to design for inclusion from the beginning.

As A4I enters its next phase, the opportunity before us is to convert these foundations into reusable platforms, transparent governance practices, and Digital Public Goods that can scale responsibly. A4I is proud to champion this effort and contribute to a model of AI innovation grounded in evidence, public value, and inclusion.

PART I

Building the Foundation

The first year translated intent into operating practice: partner deployments, field evidence loops, and early public-good capability.



From Vision to Reality

A4I's founding journey

A4I began with a belief that inclusion-focused AI must be tested where the constraints are real. Public systems operate across languages, uneven connectivity, and changing policy demands, while frontline workers already carry heavy responsibilities. The initiative was designed to bring research, engineering, implementation, and evidence into one sustained operating model.

Year 1 validated the need for this model. Each domain surfaced a different form of complexity: teachers needed support aligned with curriculum and classroom practice; ASHAs needed timely, private, and practical guidance; educators of visually impaired learners needed technology that respected tactile-first pedagogy and device realities.

What has emerged is a credible foundation. A4I established working relationships, design routines, field evidence loops, and shared engineering capabilities. It showed that responsible scale requires much more than a model, an app, or a dashboard. It requires social embedding, operational discipline, and continuous learning.

A4I moved AI innovation out of the pilot zone and into the discipline of field use, evidence, and responsible scale.

From Pilot to Scale

Building an AI impact layer

The path from pilot to scale is rarely a straight line. It is a translation journey: from technical possibility to operational reliability that fits institutional workflows.

Across the three partner programs, scale meant different things. For Sikshana Foundation, it meant helping teachers prepare more efficiently and engage classrooms better. For Vision Empower, it meant improving readiness for assistive technology and accessible content workflows. For Khushi Baby, it meant enabling ASHAs to seek reliable guidance through a familiar and private channel.

A4I's role was to build the missing middle layer: reusable engineering capability, evidence-led product evolution, implementation playbooks, and early governance patterns.

The result is a distinctive A4I loop: co-design in real settings, deploy at a meaningful scale, learn from usage, and feed those learnings back into the product and governance. This is how A4I keeps solutions grounded while still aiming for DPG-grade replication.

Year 1 was therefore a proof of execution and a proof of learning. It demonstrated reach and impact across multiple domains, while also surfacing the hard constraints that must be solved before AI systems can become dependable Digital Public Goods.

The A4I Ecosystem

Partners around one mission

A4I was structured as a collaborative ecosystem. This design matters because inclusive AI needs multiple forms of legitimacy: research credibility, engineering reliability, field trust, implementation capacity, and public-system readiness.

In this model, the partnership is not a sponsorship network. It is a production system. Microsoft brings proven innovations, funding support, and global research depth. IIIT-B brings DPG experience, public policy grounding, and the capability to engineer for scale and governance. Domain partners bring the real "operating context", the workflows, constraints, language, and trust that make adoption possible. Government support is implicit but decisive, because these solutions ride on public systems.

Institutional Partners

IIIT-Bangalore	iiit-b	Institutional home for A4I, Anchoring engineering, research, program governance, and public-good orientation.
Microsoft India	Microsoft	Anchor funding and enabling partner for inclusive AI innovation and scale. Research and technology partner supporting the original innovations, technical guidance, and scientific foundations.

Our Domain Partners

Sikshana Foundation	Sikshana FOUNDATION Love of learning	Domain partner for Shiksha copilot, supporting government schoolteachers through AI-assisted lesson planning and classroom resources.
Vision Empower	vision empower	Domain partner for accessible education, strengthening assistive technology adoption and SEEDS readiness for visually impaired learners and educators.
Khushi Baby	KhushiBaby	Domain partner for ASHA Saheli (also called ASHABot), supporting frontline health workers through WhatsApp-based AI knowledge assistance.

The A4I ecosystem combines institutional credibility, technological depth, partner trust, and field execution into a single operating model.

The A4I Team

The Steering Committee

The A4I Enablers @IIITB

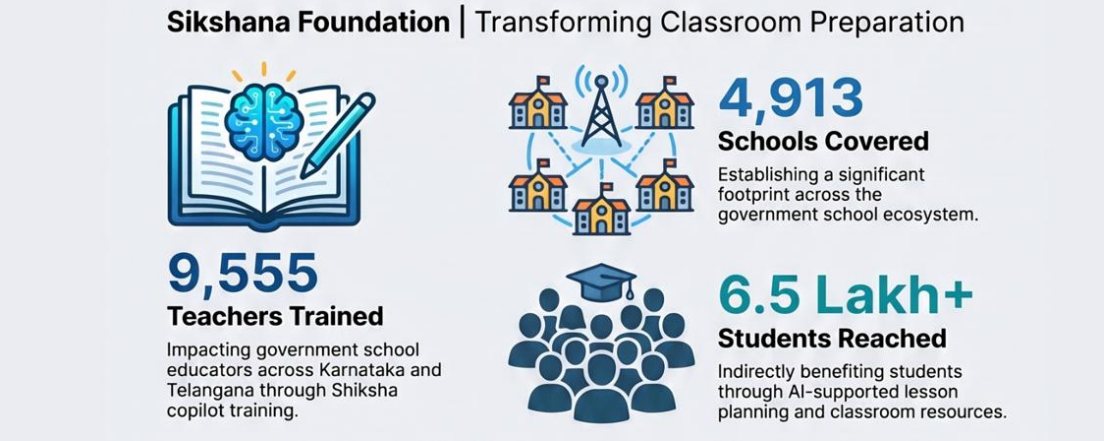
The A4I Enablers @Microsoft

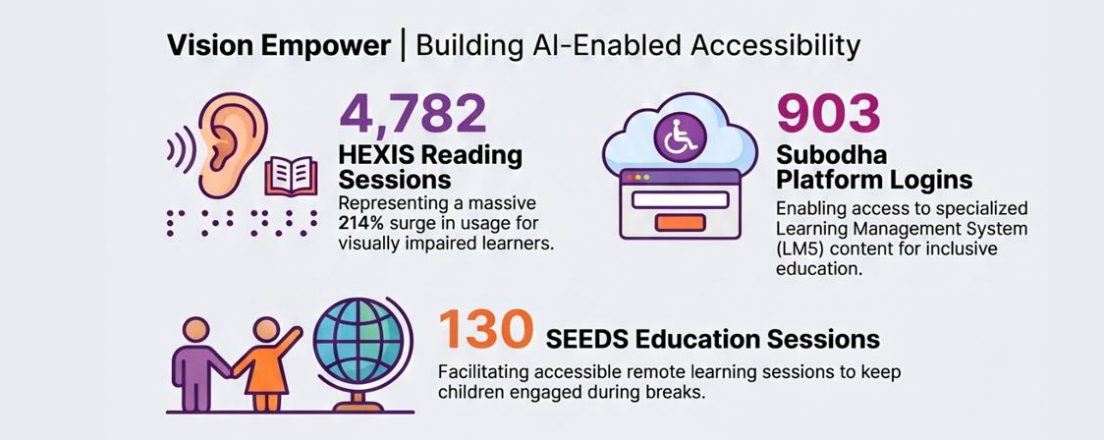
The A4I Domain Partners

Year One at a Glance

Grounded reach and adoption signals

The numbers below present Year 1 as a foundation story. They combine reach, adoption, usage, readiness, and evidence signals across the portfolio. Together, they show that A4I moved beyond concept validation into field use, partner capability building, and platform readiness across three social domains.







The numbers show a foundation year with real reach, active field use, and reusable capability emerging across education, health, and accessibility.

Journey so far - in Pictures

This section is a visual pause before the domain stories. The human reality behind the numbers:

A4I Launch



India AI Pre-Summit (hosted by Microsoft & IIITB)



ASHA engagements – ASHABot



Empowering Teachers - Shiksha copilot



Research Team in Action in Rajasthan



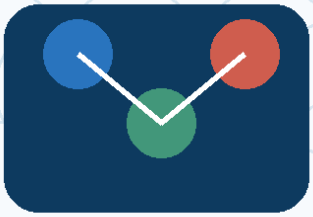
Vision Empower Accessibility Engagements



PART II

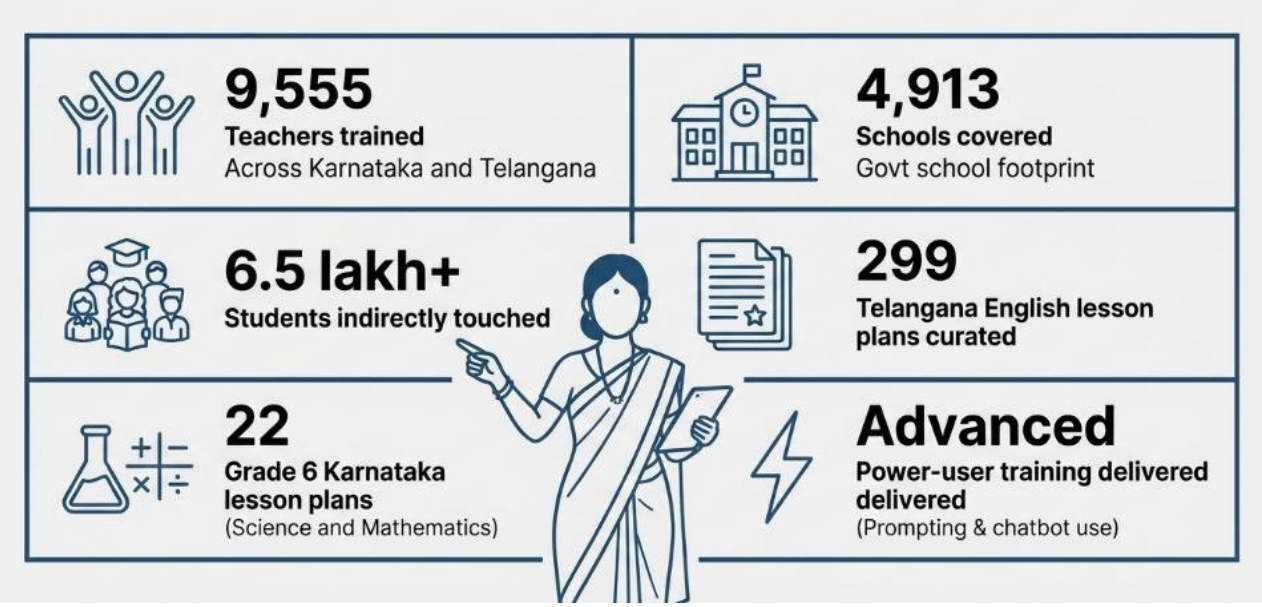
Stories of Impact

The domain stories show where A4I became real: in classrooms, special schools, and frontline health work.



Sikshana Foundation

Transforming classroom preparation through Shiksha copilot



For government school teachers, classroom instruction is only one part of the workload. Lesson preparation, assessment creation, documentation, and shifting administrative requirements all compete for time. Shiksha copilot addressed a practical pressure point: helping teachers create curriculum-aligned resources faster, while keeping the teacher in control of classroom adaptation.

In Year 1, Sikshana Foundation took Shiksha copilot into a large teacher-facing deployment across Karnataka and Telangana. The program reached 9,555 government school teachers across 4,913 schools. This exceeded the overall onboarding target and indirectly touched more than 6.5 lakh students through the teachers trained under the initiative.

The partner report highlights that Shiksha copilot supported Grades 5 to 10 across English, Science, Mathematics, and Social Science. It was used for lesson planning, question paper generation, deeper lesson adaptation, supplementary content access through lesson bots, and professional development through an AI-powered chatbot.

Impact signals from teacher practice

The most persuasive evidence from Year 1 is not the training count alone; it is the way teachers began using AI-supported resources inside real classrooms. In Belagavi, a Social Science teacher used Shiksha Copilot not only to prepare a lesson plan, but to convert an abstract geography topic into a more active classroom experience. The example is developed in the vignette below, but its significance is simple: the tool helped move lesson preparation closer to classroom engagement.

In Telangana, a Mathematics teacher reported that lesson planning, which earlier took four to five hours, could be completed in about five minutes after training. Another teacher at a Kasturba Gandhi Balika Vidyalaya used the platform to create differentiated question papers for beginner, intermediate, and advanced learners, directly addressing the reality of mixed-level classrooms.

These examples are important because they show a practical shift in teacher practice: less time spent preparing from scratch, more structured classroom activity, and more room to adapt resources to different student levels. They also show where Shiksha Copilot's value is strongest: helping teachers move from content generation to classroom use.

Implementation constraints and lessons

Year 1 also revealed that AI tools in education must keep pace with policy and curriculum shifts. In Karnataka, training was delayed by approval timelines and weekend-only training mandates. In Telangana, content customization required alignment with SCERT academic frameworks, expert availability, and translation quality assurance. These constraints are not peripheral; they are part of the adoption environment.

The core lesson is that teacher-facing AI needs more than good content generation. It must remain synchronized with curriculum changes, assessment policies, state workflows, and teacher support systems. This is why the next phase needs to move from a standalone tool toward journey-level teacher support and leader dashboards that help academic systems see adoption patterns without turning teacher support into surveillance.

Shiksha copilot showed that teacher-facing AI creates value when it saves preparation time and translates into classroom activity, not merely when it generates content.

Vignettes of grassroots impact of Shiksha copilot



Ramesh H J: Math & Science (Grades IV-VII)

Ramesh H J teaches Math and Science for grades IV to VII at CHPS Sadarahalli, Tumkur district, Karnataka. He became a power user on Shiksha copilot and has explored all its features, including generating lesson plans and using the chatbot feature for generating additional questions.

Hands-on Perimeter and Area Activity

In November, while teaching Grade 7 students, he utilized a lesson plan from Shiksha copilot for the chapter on Perimeter and Area.

To introduce the topic and engage students, he used materials like like chart paper and a compass and grouped the students. Each group was assigned different shapes, which students had to cut out from the chart paper, then measure and mark the perimeter and area of their respective shapes. This hands-on activity ensured students were actively engaged before the topic was formally completed.



Peer Advocacy & Feedback

During advanced training in February, he shared detailed support for the platform, encouraging fellow teachers and providing feedback for future enhancements.



R N Biradapatil: Social Science (Grades VI-X)

49 lesson plans generated

A teacher at Adarsha Vidyalaya RMSA, Katkol, Belagavi district, Karnataka, she has generated 49 lesson plans across multiple grades and utilizes MCQs and long-answer questions for regular student practice.

Visualizing the Lithosphere with Puzzles

While teaching a topic on the Lithosphere for Grade 8 students, she generated the lesson resource plan for the unit and engaged students in a group activity by dividing them into small groups and giving them puzzle pieces. The students had to quickly assemble the puzzle pieces, placing the Earth's layers in the correct order. Through this activity, she found that students were able to recail and understand the Earth's layers more effectively.

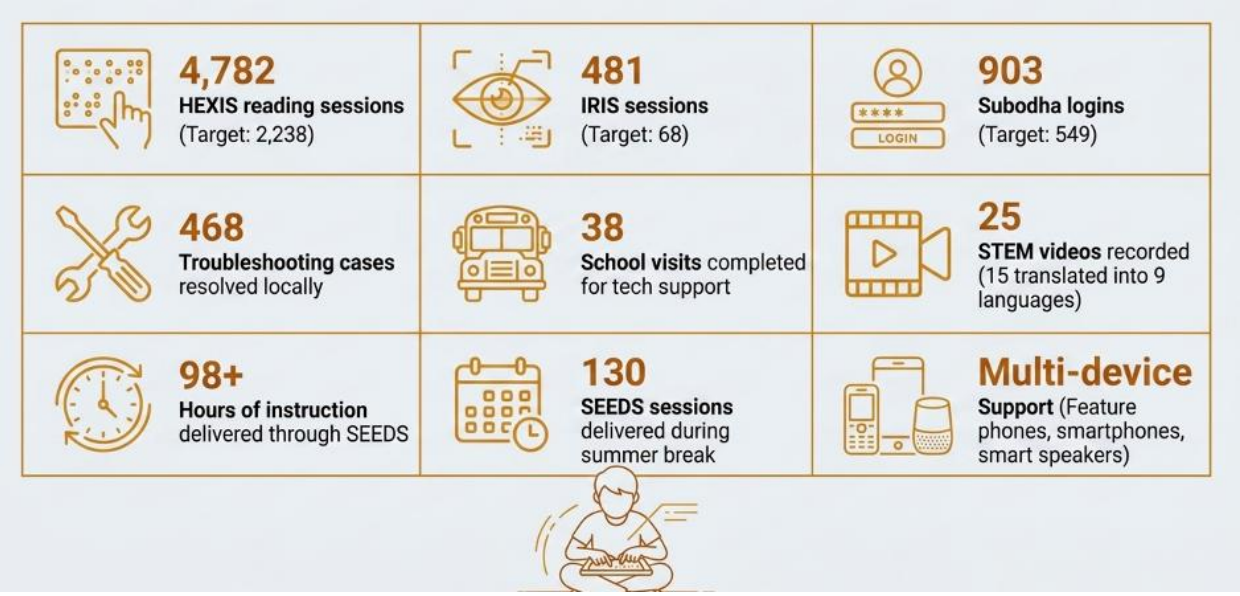


Organized & Effective Engagement

Overall, R N Biradapatil believes that the content on Shiksha copilot is very well organized and that the suggested activities are helping her to engage students more effectively.

Vision Empower

Building readiness for AI-enabled accessible education



Vision Empower's Year 1 story is not simply about deploying devices or platforms. It is about building the confidence, support systems, content pathways, and maintenance discipline that make assistive technology usable for teachers and learners with visual impairment.

The A4I project, therefore, focused on four connected forms of readiness: student readiness, teacher readiness, technology readiness, and ecosystem readiness. This framing is important because accessible education cannot be scaled through technology alone. It depends on whether students can use the tools with confidence, teachers can integrate them into tactile-first pedagogy, and schools can rely on devices and support systems when problems occur.

Vision Empower also used Year 1 to strengthen the foundation for SEEDS, envisaged as an AI-powered enabler for technology-led inclusive education. Working with the A4I Engineering team, the SEEDS infrastructure was revived, enhanced, and tested through dry runs and pilot implementations.

Impact signals from adoption and readiness

The strongest Year 1 signal from Vision Empower was movement from exposure to use. HEXIS reading sessions more than doubled the target, while IRIS sessions increased sharply, suggesting that assistive devices were not merely demonstrated but increasingly integrated into learning routines.

The partner report records the human side of this shift. Students began asking for reading materials of their choice, some used assistive devices to complete homework, and teachers reported regular sessions with students using AT devices, including news reading during school assembly. These are small but meaningful signs of independence, confidence, and normalization of accessible technology inside schools.

Content and ecosystem strengthening

Year 1 also created important content pathways. In partnership with NIEPVD and Microsoft, STEM classroom lectures were recorded at the Model School for the Visually Impaired. Twenty-five videos were recorded across Science, Mathematics, and Computer subjects, with 15 videos translated into nine languages using Azure Speech Translation. Vision Empower also mapped 13 Wipro Earthian courses and cyber safety content across staging and production environments.

This work matters because accessibility at scale depends on both content availability and distribution pathways. SEEDS can become more than a platform if it is supported by credible content intake, multi-modal transformation, accessibility checks, human review, and delivery through channels that work in real school environments.

Implementation constraints and lessons

There are operational constraints. Device maintenance is time-consuming, rural logistics make repairs difficult, translation workflows face bottlenecks, and bulk translation tools have limitations. These are not small issues. For visually impaired learners, a broken device or delayed content pathway can interrupt learning continuity.

The key lesson is that accessible education requires durable technology, local repair capacity, confidence-building support, and content workflows that respect tactile-first and audio-supported pedagogy. Year 1 established the readiness base for introducing more agentic AI-based solutions in Year 2, but the scale journey must remain grounded in maintenance, trust, and teacher agency.

Vision Empower showed that accessible AI must strengthen tactile-first learning and the support ecosystem around devices, teachers, and content.

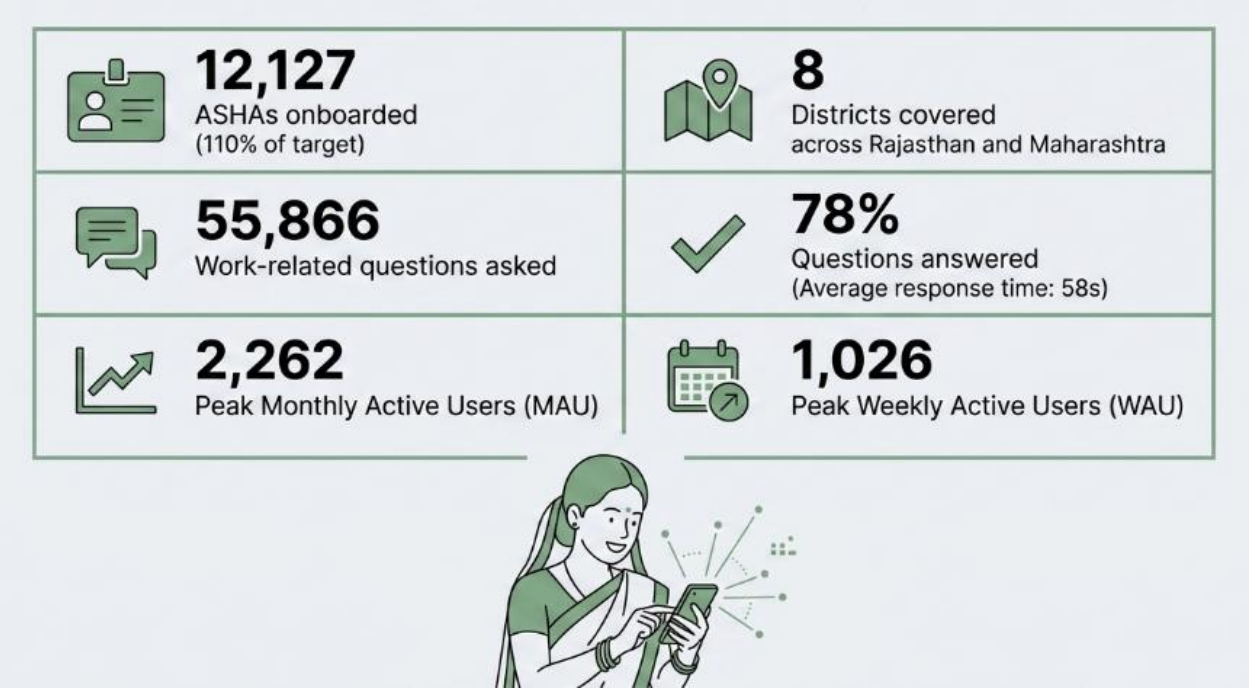
Vignettes of grassroots impact of SEEDS by Vision Empower





Khushi Baby

Supporting frontline health workers through ASHA Saheli



ASHA Saheli, also known as ASHABot, was designed around a real frontline problem: ASHAs often need timely, reliable, and context-specific information while working in the field, but supervisors are not always available, and printed or digital training materials are not easy to retrieve at the point of need.

Developed by Khushi Baby in collaboration with Microsoft Research India, ASHA Saheli used content grounded in national and state government guidelines, training modules, and resources developed by public health partners. It was implemented from March 2025 to May 2026 across eight districts in Rajasthan and Maharashtra.

The impact of Year 1 lies in the combination of reach, engagement, and learning. The deployment exceeded onboarding targets, generated a large body of real ASHA questions, and created a clear roadmap for making frontline AI more grounded, risk-aware, and useful at scale.

Impact signals from frontline use

The strongest impact signal from ASHA Saheli is that ASHAs used it for real work questions, not only for training trials. Maternal and child health dominated usage, reflecting the core of ASHA responsibilities, but the platform also received questions on supplies, payments, incentives, training, administrative processes, and practical field situations.

The implementation also showed how adoption travels through trust. WhatsApp reduced access friction, voice messages made the tool more usable for ASHAs with lower typing comfort, and weekly nudges, quizzes, and leaderboards helped sustain engagement after onboarding.

Champion ASHAs became an important adoption mechanism. Their peer influence helped normalize the tool in ways that formal training alone could not. District and block health authority support further strengthened legitimacy, allowing the solution to enter PHC and sector meetings rather than remain an external technology experiment.

IDK (I Don’t Know) analysis and risk-sensitive learning

ASHA Saheli has been trained to provide an "I don’t know” (IDK) response when a grounded, fact-based response is not available. The IDK analysis is one of the most important Year 1 contributions because it converted failure points into a roadmap for improvement. An expert review of 4,845 unanswered queries found that more than 90 percent of IDK responses across categories were due to domain knowledge gaps rather than query misinterpretation. In simple terms, the system often understood what ASHAs were asking, but the knowledge base did not yet contain enough relevant material to answer.

The largest knowledge gaps were in clinical domains, which accounted for 41 percent of IDKs, followed by operational and workflow support at 32 percent. This shows that ASHAs need both clinical guidance and day-to-day program support. A chatbot limited to static guidelines will not be enough unless it also reflects the operational reality of frontline work.

The risk analysis also sharpened the safety agenda. Maternal care had the largest number of high and very high-risk IDK queries, with 326 such queries. Child health had 106 high and very high-risk queries. These findings make a strong case for risk-aware escalation, improved danger-sign recognition, clearer referral guidance, and expert-reviewed knowledge base expansion before larger-scale roll out.

Implementation constraints and lessons

The project surfaced several practical constraints: varying digital literacy among ASHAs, lower engagement among experienced ASHAs unless advanced use cases were offered, response time above target, translation quality challenges, and difficulty operationalizing expert-in-the-loop escalation because ANMs had competing responsibilities.

The Year 1 lesson is not that AI can replace supervisory structures. It is that a well-designed AI assistant can reduce routine knowledge gaps, provide private and timely guidance, and strengthen ASHA confidence, provided it is voice-first, locally relevant, risk-aware, continuously monitored, and embedded in a trusted implementation ecosystem.

ASHA Saheli showed that frontline AI must be voice-first, risk-aware, locally relevant, and trusted through peers and public health institutions.

Vignettes of grassroots impact of ASHABot



CLINICAL & NUTRITIONAL GUIDANCE

"Addressing Anemia with Precise Nutritional Advice"

ASHAbot is helping close the gap and accelerate progress towards Anemia Muht Bharat.

During a routine counseling session, ASHA Bharati from the Jaipur district of Rajasthan sought a deeper explanation of haemoglobin deficiency to advise a worried mother. She asked ASHAbot and received a clear, step-by-step response that explained not only the target haemoglobin levels, but also practical dietary interventions—highlighting the role of leefy vegetables, pulses (dal), and iron-rich meal combinations to improve haemoglobin.

Bharti reflected, **"It's a detailed explanation, so I understand it well."**

CLINICAL & NUTRITIONAL GUIDANCE

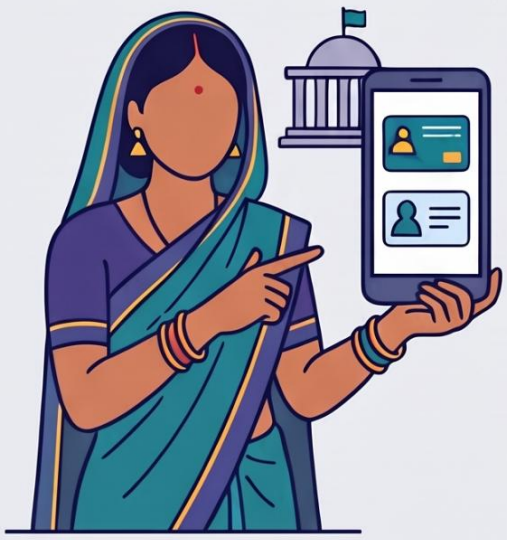
"Real-Time Support for Infant Growth Monitoring"

"How heavy should this baby be at six weeks?" -

During a home visit in Pali district of Rajasthan, ASHA worker Mani Devi was worried about an underweight infant but could not judge the severity on the spot. She opened WhatsApp, asked ASHAbot for the ideal weight for that age, and got a precise range within seconds. She followed up with, "What should I do?" and received counsel to increase the feeding frequency, along with calm, step-by-step guidance for the mother.

Mani Devi later described this as one of several interactions that changed her practice, turning hesitation into immediate, grounded action.





NAVIGATING NATIONAL HEALTHCARE BENEFITS

"Instant Guidance for Government Program Enrollment"

ASHAs are increasingly able to enroll beneficiaries into national programs at the point of care.

Kiran Kunwar, an ASHA at the Udaipur district of Rajasthan, uses ASHAbet during field visits to answer questions about Ayushman cards and ABHA IDs instantly. When community community members ask where an Ayushman card can be used, or what an ABHA ID does, Kiran consults the bot and provides immediate, accurate guidance — helping families register and access benefits faster.

"In situations like these, ASHAbot gives us a lot of information," she says, adding that the tool has made her interactions with the community more confident and efficient.

PART III

Behind the Impact

Engineering and research formed the engine room: making solutions dependable, explainable, reusable, and responsive to field realities.



A4I Engineering

Making AI dependable, reusable, and ready for responsible scale

A4I Engineering's Year 1 contribution was not limited to building features or supporting deployments. Its deeper role was to make AI solutions work in real field conditions, where content is scattered, users work across languages, connectivity is uneven, and trust can be lost quickly if systems fail.

Across ASHABot, Shiksha Copilot, and SEEDS, the team strengthened live solutions while also building reusable capabilities for future deployments. This is central to A4I's Digital Public Good direction: solutions must not only scale, but they must also be dependable, adaptable, and easier for others to adopt.

The Year 1 engineering story is best understood through two lenses: building differentiators for best-in-class solutions, and making engineering bets that support scalability and adoption.

Building Differentiators for Best-in-Class Solutions

A4I's ambition is to build solutions that are not merely functional, but among the best in their category for social-sector use. This requires strong knowledge systems, reusable content infrastructure, and adoption pathways that make mature solutions easier for others to use.

Omni-Ingest: Any content format, ready for AI

Social-sector AI depends heavily on domain content, but this content rarely arrives in an AI-ready form. It may exist as PDFs, Word documents, scanned material, audio lessons, web pages, training manuals, guidelines, or classroom resources.

Omni-Ingest creates a reusable pathway for turning this scattered content into structured, searchable, AI-ready knowledge. It helps clean, organize, enrich, and prepare content so that AI applications can retrieve grounded information more reliably. The strategic value is clear: a curriculum document for teachers, a health guideline for ASHAs, and an accessible audio resource for visually impaired learners are different in domain terms, but they share the same engineering need. Omni-Ingest reduces reinvention and gives A4I a reusable foundation for future AI solutions across domains.

Best-in-class knowledge systems: improving ASHABot

A4I took a significant step forward in strengthening the knowledge base behind ASHABot. While the solution was already grounded in national and state-level guidelines, field usage surfaced many real-world scenarios where the system did not yet have enough validated knowledge to answer confidently.

A4I worked with credible partners and medical experts to expand and refine the knowledge base. This helped reduce avoidable "I don't know" scenarios significantly. Thousands of other answers were also reviewed and improved for greater accuracy, clarity, and reliability.

This initiative is central to A4I's commitment to reliability and safety. With these improvements, ASHABot now has one of the strongest grounded knowledge bases for frontline health AI use in India, especially among similar health bots designed for field workers. The work also creates a disciplined improvement loop: real ASHA questions reveal

gaps, experts strengthen the knowledge base, and the solution becomes safer and more useful over time.

Zero Friction DPG Onboarding: Making mature solutions easier to adopt

ASHABot reached a level of maturity where it could be released as an open-source solution. This is an important milestone, but A4I's DPG ambition cannot stop at making code available. A solution becomes meaningfully reusable only when another team can understand it, configure it, deploy it, and operate it with reasonable effort.

Zero Friction DPG Onboarding is designed to make this possible. It reduces the technical burden for adopters by simplifying setup, configuration, documentation, and deployment steps. ASHABot is the first proof point for this direction. The same approach can be followed for other A4I solutions as they mature, making them not only scalable, but also easier to adopt and sustain

Key Engineering Bets Towards Scalability and Adoption

A4I Engineering also made a set of portfolio-level bets that go beyond individual solution improvements. These bets create shared capabilities that can support multiple domains, reduce reinvention, and make future deployments safer and faster.

Risk aware engineering

A4I is building safety into the engineering process, not leaving it only for post-deployment review. This means paying attention to risks related to data, accessibility, answer quality, and field reliability before they affect teachers, ASHAs, or visually impaired learners. For public-good AI, this is critical because trust can be damaged quickly if avoidable risks reach users.

Data protected, costs controlled, terminology preserved

As AI systems scale, the challenge shifts from whether the tool works to whether it can be governed responsibly. A4I is building towards infrastructure where sensitive data is better protected, usage costs remain predictable, and important health and education terms retain their meaning across languages. This is essential for responsible scale across public-service contexts.

Feature enhancement and ruggedization

Year 1 also involved significant strengthening of ASHABot, Shiksha Copilot, and SEEDS. Shiksha Copilot improved as a classroom preparation platform through stronger curriculum handling, retrieval, multi-step generation, and richer teacher-facing outputs. ASHABot became stronger through better retrieval, query handling, language support, voice interaction, and expert-reviewed knowledge workflows. SEEDS strengthened audio-first access for visually impaired learners through IVR, live audio sessions, teacher-led delivery, multilingual support, and deployment patterns suited to students without smartphones.

Across the portfolio, these engineering bets created reusable patterns: stronger knowledge pipelines, multilingual and voice-first access, domain-aware retrieval, risk-sensitive design, and solution packaging for future adoption.

A4I Engineering turns field-facing innovation into dependable infrastructure: reusable, multilingual, observable, and responsible.

A4I Research

Evidence that makes the scale responsible

A4I Research addressed a central question for responsible AI scale: what does it take for AI-enabled solutions to become sustainable, inclusive, and socially embedded Digital Public Goods?

In Year 1, the research team studied ASHABot, Shiksha Copilot, and Vision Empower through fieldwork, interviews, observations, and qualitative analysis across health, education, and accessible education contexts. Its contribution was not limited to evaluation. It created a field-based intelligence layer for A4I, helping explain why adoption happens, why it stalls, and what design, governance, and implementation choices are needed when AI enters real public-sector workflows.

The core research contribution

The key intellectual contribution of Year 1 was the development of a Scaling AI as Digital Public Goods framework. The framework connects technical readiness, such as openness, standards, documentation, and safeguards, with socio-technical readiness, such as user agency, institutional support, local adaptation, knowledge transfer, and trusted networks.

This matters because scale is not just more users or more geographies. In A4I's context, scale means adoption breadth, workflow depth, institutional fit, technical robustness, ethical oversight, and the ability of local actors to adapt and sustain the system.

What the research discovered

Area	Research insight	Why it matters
ASHABot and network effects	ASHAs are not only service providers. They are knowledge brokers within trusted community and peer networks.	Health AI must strengthen peer learning, voice-first access, workflow fit, and risk-aware escalation rather than bypassing existing networks.
Shiksha Copilot and institutional logics	Adoption was shaped by the alignment between government policy, teacher professionalism, and implementation management.	Education AI must stay aligned with curriculum, assessment rules, teacher workload, and academic leadership systems.
Vision Empower and accessible education	Visually impaired education is tactile-first, individualized, and dependent on teacher judgment.	Accessible AI must support independence, durable devices, local support, and teacher-mediated learning rather than automate instruction.
Scaling AI as DPGs	Technical readiness alone is not enough for responsible scale.	A4I must combine reusable infrastructure with institutional fit, ethical safeguards, local agency, and support ecosystems.

How research influenced the program

Research findings were shared with Microsoft Research, A4I Engineering, and domain partners, creating a feedback loop from field evidence to program design.

For ASHA Saheli, research and IDK analysis strengthened the case for knowledge-base expansion, risk detection, voice-first workflows, and practical operational support. For Shiksha Copilot, research highlighted the importance of alignment with policy shifts, teacher workflows, and leader dashboards. For Vision Empower, it reinforced the need to build assistive technology adoption around tactile-first pedagogy, device reliability, local support, and teacher confidence.

The strategic value of A4I Research is that it prevents the initiative from mistaking deployment for impact. It replaces assumptions with field evidence and makes responsible scale a design discipline.

Research outputs and visibility

Year 1 research produced strong academic and policy visibility: three full papers, three accepted abstracts, and one policy note and toolkit on AI scale-up and governance. These outputs positioned A4I not only as an implementation initiative, but also as a contributor to the emerging knowledge base on responsible AI and Digital Public Goods.

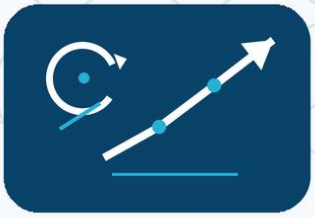
Output type	Title / focus	Conference or forum	Status
Full paper	Scaling Technology as Digital Public Goods: A Conceptual Framework	2nd India Conference on Information Systems, InCIS 2026	Accepted
Full paper	Frontline Workers, Chatbots and Care: Social and Technical Factors Shaping Network Effects in AI for Maternal Health	19th IFIP WG 9.4 Conference	Accepted
Full paper	Scaling Generative AI in Public School Education: An Institutional Logics Perspective from the Global South	19th IFIP WG 9.4 Conference	Accepted
Abstract	Cyborg Teaching in Splintered Schools: A qualitative study on the use of an AI-EdTech tool in India's public education system	6th India Public Policy Network Conference	Accepted
Abstract	Between Design and Practice: AI, Non-Use, and Lived Policy in ASHAs' Digital Work	6th India Public Policy Network Conference	Accepted
Abstract	Frontline Workers, Chatbots and Healthcare: Honouring social contracts to harness AI use in community health	Bringing Evidence to Public Health Policy, EPHP	Accepted
Policy output	Policy note and toolkit for AI scale-up and governance	A4I dissemination and partner use	Completed

A4I Research acts as deployment intelligence, replacing assumptions with field evidence and making responsible scale a design discipline.

PART IV

Reflections and the Road Ahead

The closing section converts Year 1 learning into a disciplined path for journey-scale, responsible AI platforms.



What We Learned

Five grounded lessons from Year 1

Field value is created when AI fits the workflow.

Across teachers, ASHAs, and educators of visually impaired learners, adoption was strongest when AI reduced a real work burden. Shiksha Copilot helped teachers move faster from preparation to classroom activity. ASHA Saheli helped frontline workers seek timely guidance in a familiar channel. SEEDS and assistive technologies strengthened learning continuity where smartphones, content access, and specialized support were limited. The lesson is clear: AI creates value when it fits the user's day, not when it simply demonstrates technical capability.

Scale depends on the surrounding ecosystem.

Year 1 showed that adoption is shaped by much more than the product. Government approvals, curriculum alignment, expert availability, translation quality, device maintenance, training schedules, peer champions, and local supervisors all affected the pace and depth of use. This is not a weakness of the model. It is the reality of public-service scale. A4I's role is to make these ecosystem dependencies visible and design around them.

Reliability, safety, and trust must be engineered in.

In social-sector AI, trust can be lost quickly if a system gives weak answers, fails during use, mistranslates critical terms, or cannot be recovered easily. Year 1 reinforced the importance of rugged platforms, stronger knowledge bases, risk-aware answer improvement, multilingual reliability, and observable systems. Engineering is therefore not a support function alone. It is one of the foundations of responsible impact.

Inclusion requires deliberate design choices.

Accessible and equitable use does not happen automatically. Voice-first interaction, tactile-first pedagogy, local language support, phone-based access, device reliability, and teacher-mediated learning all mattered in Year 1. The most inclusive solutions were those that respected real constraints instead of assuming ideal digital access. This principle must guide every future A4I solution from the start.

Reusable infrastructure turns learning into scale.

Year 1 created more than partner-specific deployments. It produced reusable patterns: content ingestion, grounded retrieval, knowledge-base improvement, multilingual support, zero-friction DPG onboarding, research feedback loops, and governance-oriented design. These assets reduce reinvention and make future deployments faster, safer, and easier to adapt. This is how A4I can move from successful implementations to a public-good platform.

→ The Path Ahead

From AI assistance to journey-scale public-good platforms

Year 2 can build on a stronger foundation created in Year 1. The opportunity is to move from AI assistance toward journey-scale platforms that connect planning, action, feedback, escalation, and measurement, while staying grounded in the operational realities surfaced during the first year.

In education, this means moving from lesson planning support toward teacher journeys that connect classroom preparation, teaching support, assessments, student gaps, and leader dashboards. In accessible education, it means strengthening SEEDS as a multi-modal content journey while stabilizing the assistive technology ecosystem that makes learning possible. In healthcare, it means evolving ASHA Saheli from an information chatbot toward a risk-aware, voice-first, workflow-aligned assistant with stronger knowledge coverage and escalation logic.

Across all domains, A4I's role is to convert these learnings into reusable assets: reference architectures, multilingual infrastructure, governance controls, evidence frameworks, implementation playbooks, and partner enablement models. This is the path from promising AI tools to Digital Public Goods that can be responsibly adapted and scaled.

The next phase is about moving from AI assistance to governed journeys that connect action, feedback, escalation, and measurable public value.

Closing Note

A foundation year for responsible AI at scale

A4I's first year showed that AI for social impact must be built in the field, not merely deployed there. The work across education, accessibility, and health demonstrated practical reach and early adoption, but its deeper contribution was the creation of a learning system: partners implementing, engineering strengthening, research explaining, and governance preparing the path to scale.

The next chapter is to convert this foundation into dependable public-good infrastructure: grounded in evidence, strengthened by partners, and designed for the people who hold India's last-mile systems together.

The promise of A4I is not AI by itself. It is AI as a shared public capability: trusted in the field, inclusive by design, reusable across contexts, and accountable to the people it is meant to serve.